arXiv:1703.08774v1 [cs.LG] 26 Mar 2017

Who Said What: Modeling Individual Labelers Improves Classification

Melody Y. Guan*
Varun Gulshan
Andrew M. Dai
Geoffrey E. Hinton
{melodyguan,varungulshan,adai,geoffhinton}@google.com
Google Brain
1600 Amphitheatre Pwky
Mountain View, California 94043

ABSTRACT

Data are often labeled by many different experts with each expert only labeling a small fraction of the data and each data point being labeled by several experts. This reduces the workload on individual experts and also gives a better estimate of the unobserved ground truth. When experts disagree, the standard approaches are to treat the majority opinion as the correct label or to model the correct label as a distribution. These approaches, however, do not make any use of potentially valuable information about which expert produced which label. To make use of this extra information, we propose modeling the experts individually and then learning averaging weights for combining them, possibly in sample-specific ways. This allows us to give more weight to more reliable experts and take advantage of the unique strengths of individual experts at classifying certain types of data. Here we show that our approach leads to improvements in computer-aided diagnosis of diabetic retinopathy. We also show that our method performs better than competing algorithms by Welinder and Perona, and by Mnih and Hinton. Our work offers an innovative approach for dealing with the myriad real-world settings that use expert opinions to define labels for training.

1 INTRODUCTION

Over the last few years, deep convolutional neural networks have led to rapid improvements in the ability of computers to classify objects in images and they are now comparable with human performance in several domains. As computers get faster and researchers develop even better techniques, neural networks will continue to improve, especially for tasks where it is possible to get a very large number of accurately labeled training examples. In the near future, we can expect neural networks to start serving as alternatives to human experts. We would, in fact, like the neural networks to perform much better than the human experts used to provide the training labels because these training labels are often unreliable as indicated by the poor agreement between different experts (55.4% for the datasets we consider) or even between an expert and the same expert looking at the same image some time later (70.7%).[1] Intuitively, we would expect the quality of the training labels to provide an upper bound on the performance of the trained net. In the first part of the paper we show that this intuition is incorrect.

The main contribution of the paper is to show that there are significantly better ways to use the opinions of multiple experts than simply treating the consensus of the experts as the correct label or using the experts to define a probability distribution over labels.

2 MOTIVATION

2.1 Beating the teacher

To demonstrate that a trained neural net can perform far better than its teacher we use the well-known MNIST benchmark for which the true labels are known and we create unreliable training labels by corrupting the true labels. This corruption is performed just once per experiment, before training starts, so the noise introduced by the corruption cannot be averaged away by training on the same example several times. MNIST has 60,000 training images and 10,000 test images of isolated, normalized, hand-written digits and the task is to classify the image into one of ten classes. Each image has 28×28 pixels. For the purposes of this demonstration, we used a very simple neural net containing two hidden convolutional layers each with 1,024 rectified linear units and 64 patches followed by a fully connected hidden layer of 32 rectified linear units followed by a 10-way softmax layer. We trained the net on 50,000 examples using stochastic gradient descent on mini-batches of size 200 with the Adam optimizer [14] and we used the remaining 10,000 training images as a validation set to select good values for the learning rate and the magnitude of the initial random weights. On the test data, the net that performed best on the validation set had an test error rate of 1.01% when the training labels were all correct.[2] If the labels are corrupted by changing each label to one of the other nine classes with a probability of 0.5, the test error rate only rises to 2.29%. Even if each training label is changed to an incorrect label with probability 0.8 so that the teacher is wrong 80% of the time, the trained net only gets 8.23% test error. If the teacher is even less reliable there comes a point at which the neural net fails to "get the point" and its error rate rises catastrophically, but this does not happen until the teacher is extremely unreliable as is shown in Figure 1.

*Work done as a member of the Google Brain Residency program (g.co/brainresidency).

[1] Inter-grader variability is a well-known issue in many settings in which human interpretation is used as a proxy for ground truth, such as radiology [7] or pathology [6].

[2] We did hyperparameter tuning for data where each training label was changed to an incorrect label with probability 0.8. If we had tuned for data where all labels were correct, the corresponding error rate would have been lower.

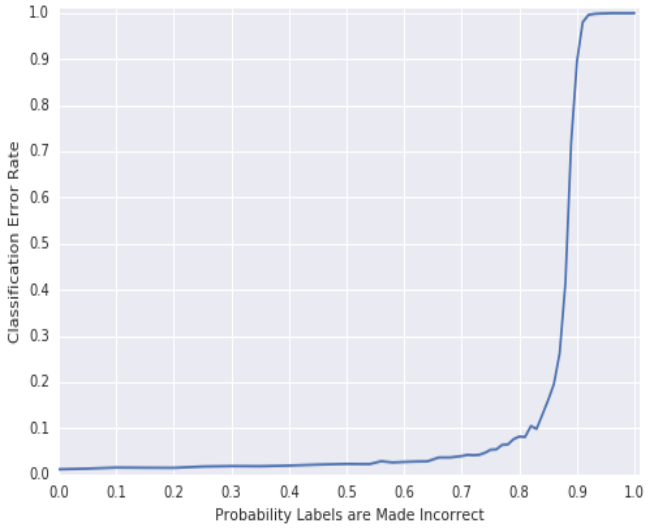

Figure 1: Performance of a deep neural net when trained with noisy labels.

This demonstration shows that the performance of a neural net is not limited by the accuracy of its teacher, provided the teacher's errors are random. One obvious question is how many noisily labeled training examples are worth a correctly labeled training example. In Appendix B we show that this question can be answered, at least approximately, by computing the mutual information between the label and the truth.

2.2 Making better use of noisy labels

We are interested in datasets of medical images where many different doctors have provided labels but each image has only been labeled by a few doctors and most of the doctors have only labeled a fairly small fraction of the images. We expect that some doctors will be more reliable than others and we would like to give more weight to their opinions. We also expect that the doctors will have received different training and may have experienced different distributions of images so that the relative reliability of two doctors may depend on both the class of the image and on properties of the image such as the type of camera it was taken with. In this paper we focus on datasets of images used for screening diabetic retinopathy because neural networks have recently achieved human-level performance on such images [10] and if we can produce even a relatively small improvement in the state-of-the-art system it will be of great value.

There is more information in the particular labels produced by particular doctors than is captured by simply averaging the opinions of all the doctors who have labeled a particular image and treating this distribution as the correct answer. The amount of constraint that a training case imposes on the weights of a neural network depends on the amount of information required to specify the desired output. So if we force the network to predict what each particular doctor would say for each particular training case we should be able to get better generalization to test data, provided this does not introduce too many extra parameters. For a K-way classification task, we can replace the single softmax [9] that is normally used by as many different K-way softmaxes as we have doctors. Of course, there will be many doctors who have not labeled a particular training image, but this is easily handled by simply not backpropagating any error from the softmaxes that are used to model those doctors. At test time we can compute the predictions of all of the modeled doctors and average them. Our belief is that forcing a neural network to model the individual doctors and then averaging at test time should give better generalization than simply training a neural network to model the average of the doctors.

We should also be able to do better than just averaging the opinions of the modeled doctors. After we have finished learning how to model all of the individual doctors we can learn how much to weight each modeled doctor's opinion in the averaging. This allows us to downweight the unreliable doctor models.

2.3 Diabetic retinopathy classification

Diabetic retinopathy (DR) is the fastest growing cause of blindness worldwide, with nearly 415 million diabetics at risk [8]. Early detection and treatment of DR can reduce the risk of blindness by 95% [12]. One of the most common ways to detect diabetic eye disease is to have a specialist examine pictures of the back of the eye called fundus images and rate them on the International Clinical Diabetic Retinopathy scale [16], defined based on the type and extent of lesions (e.g. microaneurysms, hemorrhages, hard exudates) present in the image. The image is classified into one of 5 categories consisting of (1) *No DR*, (2) *Mild NPDR (non-proliferative DR)*, (3) *Moderate NPDR*, (4) *Severe NPDR*, and (5) *Proliferative DR* (Figure 2). Another important clinical diagnosis that can be made from the fundus image is the presence of diabetic macular edema (DME). While this work focuses only on the 5 point grading of DR, the findings should be applicable to DME diagnosis as well.

Most of the prior work on diabetic retinopathy classification focuses on obtaining a single ground truth diagnosis for each image, and then using that for training and evaluation. Deep learning has recently been used within this setting by Gulshan et al. [10] who show a high sensitivity (97.5%) and specificity (93.4%) in the detection of referable DR (moderate or more severe DR).[3]

In this work we explore whether, in the context of data in which every example is labeled by multiple experts, a better model can be trained by predicting the opinions of the individual experts as opposed to collapsing the many opinions into a single one. This allows us to keep the information contained in the assignment of experts to opinions, which should be valuable because experts labelling data differ from each other in skill and area of expertise (as is the case with our opthalmologists, see Figure 3). Note that we still need a single opinion on the test set to be able to evaluate the models. To that end, we use a rigorous adjudicated reference standard for evaluation, where a committee of three retinal specialists resolved disagreements by discussion until a single consensus is achieved.

[3] In a recent Kaggle machine-learning competition [11] for DR, all the winning models also used deep learning. There have been non-deep-learning attempts as well. On the Messidor-2 dataset [4], Abràmoff et al. [2] report a sensitivity of 96.8% at a specificity of 59.4% for detecting referable DR while Solanki et al. [22] report a sensitivity of 93.8% at a specificity of 72.2%.

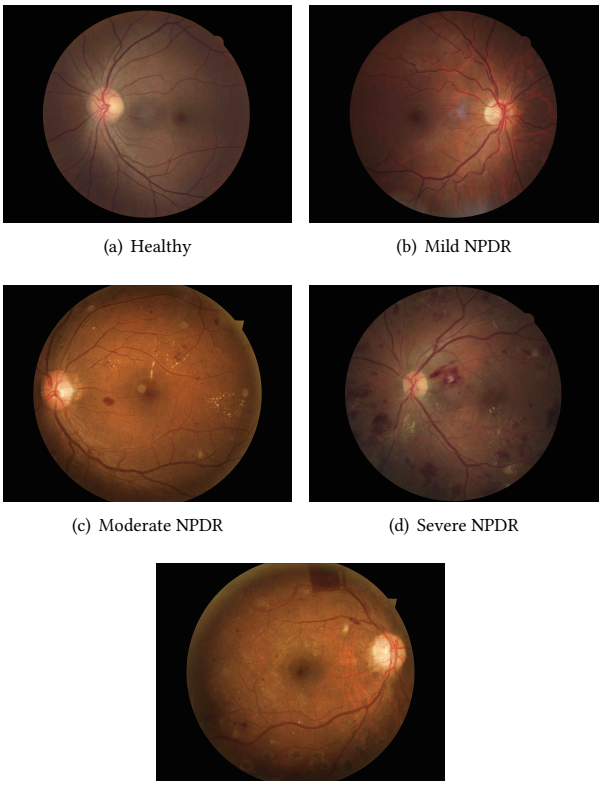

(a) Healthy (b) Mild NPDR

(c) Moderate NPDR (d) Severe NPDR

(e) Proliferative DR

Figure 2: Sample fundus images from each DR class.

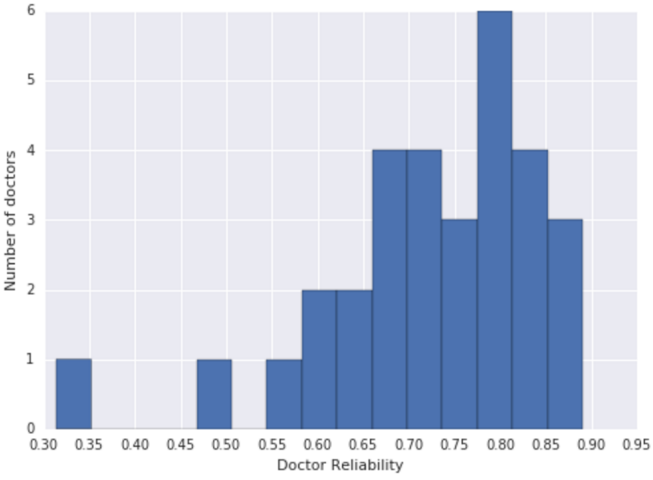

Figure 3: Histogram of doctor reliabilities. These were calculated from Welinder and Perona [24]'s expectation-maximization algorithm on our training data.

3 METHODS

3.1 Model Architecture

We considered a sequence of models of increasing complexity for training the diabetic retinopathy classifier (Figure 5). The neural network base used in this work is the Inception-v3 architecture proposed by Szegedy et al. [23].

- Baseline Net (BN): Inception-v3 trained on average opinions of doctors; a TensorFlow reimplementation of the model used in Gulshan et al. [10] (see Section 4.3 for differences).
- Doctor Net (DN): BN extended to model the opinions of each of the doctors (31 in total, see Section 4.2).
- Weighted Doctor Net (WDN): Fixed DN with averaging weights for combining the predictions of the doctor models learned on top, one weight per doctor model.
- Image-specific WDN (IWDN): WDN with averaging weights that are learned as a function of the image.
- Bottlenecked IWDN (BIWDN): IWDN with a small bottleneck layer for learning the averaging weights.

For BN, the outputs of the last hidden layer of Inception were used to compute the logits used in the 5-way softmax output layer. For DN, the opinions of each doctor were modeled using a separate softmax for each doctor, while Inception weights were shared. For evaluation, the predictions from the softmax "doctor models" were arithmetically averaged to give a single 5-class prediction. For subsequent nets, the parameters and predictions of the DN model were frozen and only the averaging weights for the doctor models were learned. For WDN, one averaging weight per doctor was trained, used across all images. For IWDN, these averaging weights were made image-dependent by letting them be a function of the last hidden layer of Inception. For BIWDN, a linear bottleneck layer of size 3 was added between the last hidden layer of Inception (which has dimension 2048) and the 31-way softmax of IWDN as a precautionary measure against model underfitting; a bottleneck layer of this size reduced the number of trainable parameters about 10 times.

Rather than directly learning the averaging weight for each doctor model (B)(I)WDN, we learned averaging logits for each model that we could then pass through a softmax to produce averaging weights that are guaranteed to be positive. To train the averaging logits, we used the opinions of the doctors who actually labeled a training image to define the target output distribution for that image (Appendix C.2 discusses an alternative target). We then combined the predictions of the models of all the other doctors using the weights defined by their current averaging logits. Finally we updated our parameters by backpropagating with the cross entropy loss between the target distribution and the weighted average prediction. This way all of the training cases that a doctor did *not* label can be used to learn the averaging logit for that doctor, and no extra data were needed beyond those used to learn the weights of DN. Moreover, if a doctor model has similar performance to other doctor models but makes very different errors it will tend to be upweighted because it will be more useful in the averaging. This upweighting of diverse doctor models would not occur if we had computed the reliabilities of the doctors separately.

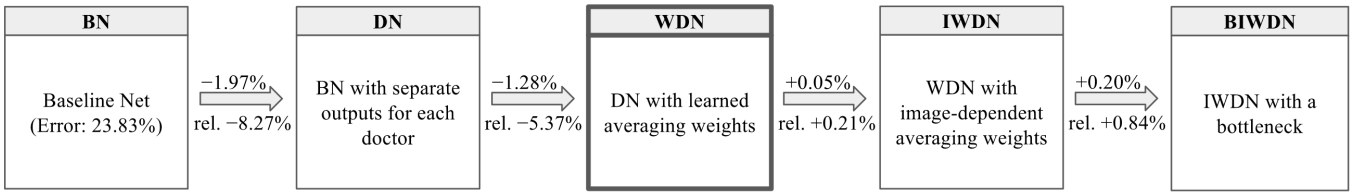

Figure 4: Description of nets used in paper. The baseline net had 5-class classification error rate of 23.83% on the test dataset. The numbers above arrows refer to absolute changes in test error while the numbers below arrows refer to relative changes (negative values represent improvements). WDN (highlighted) was the optimal net.

Table 1: Prediction inputs to cross entropy loss for each model during training. The notation is given in the text. Note that the target is always $\frac{1}{|I|}\sum_{i\in I} l_i$.

Model	Training	Evaluation
BN	$p_\emptyset$	$p_\emptyset$
DN	$p_i, \forall i \in I$	$\frac{1}{31}\sum_{i=1}^{31} p_i$
(B)(I)WDN	$\frac{\sum_{i\in I} p_i \cdot w_i}{\sum_{i\in I} w_i}$	$\sum_{i=1}^{31} p_i \cdot w_i$

For a single image, let I be the set of indices of the doctors who actually graded that image. Let the label of doctor $i \in I$ be l_i. For every doctor $j \in \{1, 2, \ldots, 31\}$, denote the prediction of its model p_j. Let $p_\emptyset$ be the prediction of the model of the average doctor in BN. For WDN, IWDN, and BIWDN, let w_j be the averaging weight for the jth modeled doctor, where $\sum_j w_j = 1$. Note that p_j is a 5-dimensional vector and w_j is a scalar. The explicit inputs of the cross entropy loss being minimized during training of each model are shown in Table 1 and post-Inception computations are shown schematically in Figure 5. In the case of DN, the cross entropy losses of the individual doctor models were summed to get the total loss for each training example.

3.2 Estimating doctor reliability with EM

Since the foundational work of Dawid and Skene [3], who model annotator accuracies with expectation-maximization (EM), and Smyth et al. [21], who integrate the opinions of many experts to infer ground truth, there has a large body of work using EM approaches to estimate accurate labels for datasets annotated by multiple experts [18, 19, 25]. Representatively, Welinder and Perona [24] use an online EM algorithm to estimate abilities of multiple noisy annotators and to determine the most likely value of the labels. We calculated updated labels by executing Welinder and Perona [24]'s method on our human doctors and used these updated labels to train BN, as a competing algorithm for our DN method. Welinder and Perona [24] also actively select which images to label and how many labels to request based on the uncertainty of their estimated ground truth values and the desired level of confidence, and they select and prioritize which annotators to use when requesting labels. We do not use these other aspects of their algorithm because labels for all images in our dataset have already been collected.

3.3 Modeling label noise

Mnih and Hinton [15] describe a deep neural network that learns to label road pixels in aerial images. The target labels are derived from road maps that represent roads using vectors. These vectors are converted to road pixels by using knowledge of the approximate width of the roads so the target labels are unreliable. To handle this label noise, Mnih and Hinton [15] propose a robust loss function that models asymmetric omission noise.

They assume that a true, unobserved label $\mathbf{m}$ is first generated from a $w_m \times w_m$ image patch $\mathbf{s}$ according to some distribution $p(\mathbf{m}|\mathbf{s})$, and the corrupted, observed label $\tilde{\mathbf{m}}$ is then generated from $\mathbf{m}$ according to a noise distribution $p(\tilde{\mathbf{m}}|\mathbf{m})$. The authors assume an asymmetric binary noise distribution $p(\tilde{m}_i|m_i)$ that is the same for all pixels i. They assume that conditioned on $\mathbf{m}$, all components of $\tilde{\mathbf{m}}$ are independent and that each $\tilde{m}_i$ is independent of all $m_{j\neq i}$. The observed label distribution is then modeled as:

$$p(\tilde{\mathbf{m}}|\mathbf{s}) = \prod_{i=1}^{w_m^2} \sum_{m_i} p(\tilde{m}_i|m_i)p(m_i|\mathbf{s})$$

We used a multi-class extension of their method on DN, modeling the noise distribution prior for all doctors d with the parameters:

$$\theta_{ll'} = p(\tilde{m}_d = l'|m_d = l)$$

where $l, l' \in \{1, 2, 3, 4, 5\}$. We estimated $\theta_{ll'}$ using the 5×5 confusion matrix between individual and average doctor opinions on training images. Treating the average doctor opinion as the true label, we converted each doctor's individual count matrix into proportions and then averaged these proportions across all doctors. We trained this model by minimizing the negative log posterior, $-log(p(\tilde{\mathbf{m}}|\mathbf{s}))$. Our variant of Mnih and Hinton [15] is an alternative way to improve upon DN to our proposed approach of learning averaging weights.

4 EXPERIMENTAL SETUP

4.1 Neural network training

The optimization algorithm used to train the network weights was distributed stochastic gradient descent (SGD) [1] with the Adam optimizer on mini-batches of size 8. We trained using TensorFlow with 32 replicas and 17 parameter servers, with 1 Tesla K80 GPU per replica. To speed up the training, we used batch normalization [13], pre-initialization of our Inception network using weights from the network trained to classify objects in the ImageNet dataset [20], and the following trick: we set the learning rate on the weight matrix producing prediction logits to one-tenth of the learning rate for

the other weights. We prevented overfitting using a combination of L1 and L2 penalties, dropout, and a confidence penalty [17], which penalizes a model for having an output distribution with low entropy. At the end of training, we used an exponentially decaying average of the recent parameters in the final model.

We tuned hyperparameters and picked model checkpoints for early stopping on the validation dataset, using 5-class classification error rate as the evaluation metric. The optimal values for these hyperparameters are displayed in Table 2 and the hyperparameter search spaces are listed in Appendix D. When evaluating on the test set we averaged the predictions for the horizontally and vertically flipped versions (4 in total) of every image.

We also trained a version of BN where the output prediction is binary instead of multi-class, as was done in Gulshan et al. [10]. The binary output was obtained by thresholding the 5-class output at the *Moderate NPDR* or above level, a commonly used threshold in clinics to define a referable eye condition. For this BN-binary network, the area under the ROC curve was used as the evaluation metric on the validation set.

To deal with differences in class distribution between the datasets (Table 3), we used log prior correction during evaluation. This entails adding to the prediction logits, for each class, the log of the ratio of the proportion of labels in that class in the evaluation dataset to the proportion of labels in that class in the training set. Our assumed test class distribution for computing the log prior correction was the mean distribution of all known images (those of the training and validation sets).[4] So for each image under evaluation we update the prediction logit for class c by adding:

$$log\left(\frac{q_{valid}(c)}{q_{train}(c)}\right) \qquad \text{for the validation dataset, and}$$

$$log\left(\frac{q_{valid\cup train}(c)}{q_{train}(c)}\right) \qquad \text{for the test dataset,}$$

where $q(c)$ is the proportion of labels in that class. We saw improvement from the application of log prior correction and all our reported results use it. See Appendix C.1 for another way we attempted to correct for differences in class distribution.

4.2 Datasets

The training dataset consists of 126, 522 images sourced from patients presenting for diabetic retinopathy screening at sites managed by 4 different clinical partners: EyePACS, Aravind Eye Care, Sankara Nethralaya, and Narayana Nethralaya. The validation dataset consists of 7,804 images obtained from EyePACS clinics. Our test dataset consists of 3,547 images from the EyePACS-1 and Messidor-2 datasets. More details on image sourcing can be found in Appendix E.

Each of the images in the training and validation datasets was graded by at least one of 54 US-licensed ophthalmologist or ophthalmology trainee in their last year of residency (postgraduate year 4). For training the doctor models, we used the 30 ophthalmologists who graded at least 1,000 images, and we lumped the remaining doctors as a single composite doctor to avoid introducing doctor-specific parameters that are constrained by fewer than 1,000 training cases. Meanwhile, the labels for the test set were obtained through an adjudication process: three retina specialists graded all images in the test dataset, and discussed any disagreements as a committee until a consensus label was obtained.

We scale normalized our images by detecting the circular fundus disk and removing the black borders around them. We used images at a resolution of 587×587 pixels and we augmented our training data with random perturbations to image brightness, saturation, hue, and contrast.

4.3 Our baseline vs published baseline

This section describes multiple ways in which our baseline differs from that of Gulshan et al. [10]. For these reasons, results from this paper's own BN should be used for model comparisons with DN, WN, IWDN, and BIWDN rather than numbers from Gulshan et al. [10].

- Unlike in Gulshan et al. [10], we remove grades of doctors who graded test set images from training and validation sets to reduce the chance that the model is overfitting on certain experts. This removal handicaps our performance vis-à-vis their paper, especially because we exclude the most expert doctors (the retinal specialists) during model development, but ensures generalizability of our results.
- We use different datasets, and in particular our adjudicated test set has gold standard labels that are meant to represent the ground truth.
- We train with 5-class loss instead of binary loss (see sections 4.1 and 5.1).
- If a doctor grades a single image multiple times, as often occurs, Gulshan et al. [10] treats these as independent diagnoses while we collapse these multiple diagnoses into a single diagnosis which may be a distribution over classes.
- We employ higher resolution images (587×587 pixels versus 299×299) and image preprocessing and theoretical techniques unused in Gulshan et al. [10] (section 4.2).

More differences are discussed in Appendix F.

5 SUMMARY OF RESULTS

We ran 10 replicates of each model and averaged the resulting metrics, which are reported in Table 5. For full comparability of models we used the same 10 replicates reported for DN to serve as the fixed part of the model for training the WDN, IWDN, and BIWDN replicates.

5.1 Training with five-class loss beats training with binary loss even on binary metrics

We found that training BN with a 5-class loss improves test binary AUC compared to training with a binary loss, as did Gulshan et al. [10], even when validating the former on 5-class training error instead of binary AUC (Table 4). Test binary AUC was raised by 1.53% from 95.58% from using the multi-class loss. Intuitively this fits with our thesis that generalization is improved by increasing the amount of information in the desired outputs. All results reported in Table 5 and subsequent sections were obtained from training with 5-class loss.

[4] We did not follow the standard machine learning paradigm of assuming that the test dataset has the same distribution as the validation dataset because our test data are sourced differently from the validation data.

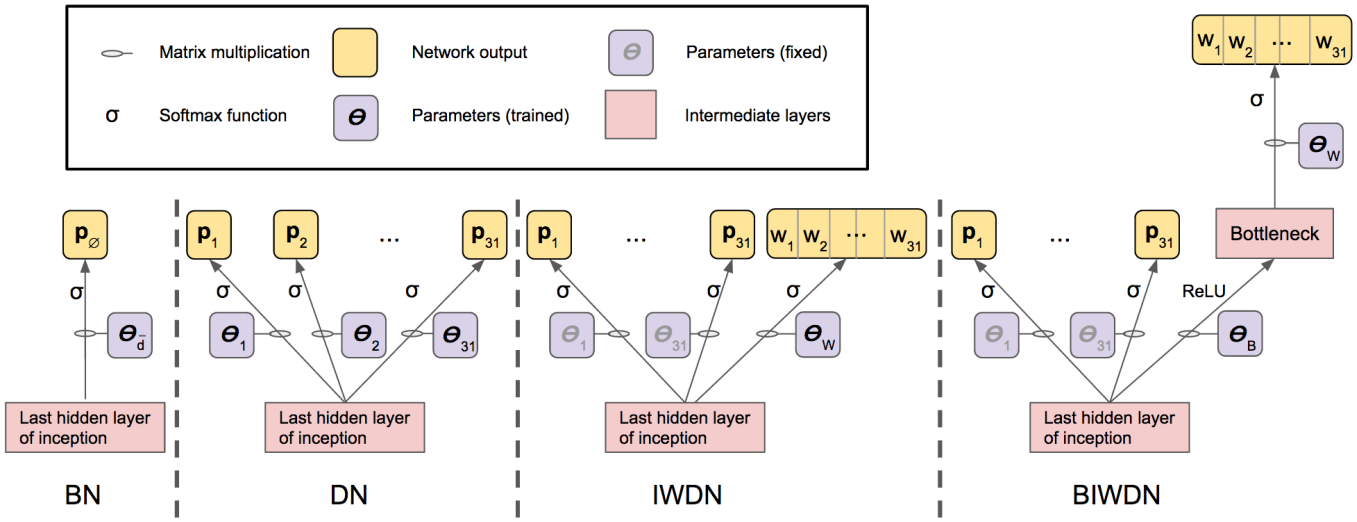

Figure 5: Schematic diagram of nets. These schematics show how the parameters, network outputs, and averaging weights for doctor models are connected. Table 1 lists how the outputs are used in a loss function for training. In WDN (not shown in figure), the averaging logits are not connected to the last hidden layer of Inception and are just initialized from a constant vector.

Table 2: Optimal Hyperparameters from Grid Search. Note that the learning rate for doctor models is one-tenth the learning rate for the rest of the network listed here.

Hyperparameter	BN binary	BN	BN EM	DN	DN mnih	WDN	IWDN	BIWDN
Learning rate	0.0001	0.0003	0.0003	0.001	0.0003	0.03	1×10^{-6}	3×10^{-7}
Dropout for Inception	0.75	0.95	0.95	0.85	0.95	-	-	-
Dropout for output heads	0.8	0.85	0.85	0.9	0.9	-	-	-
Entropy weight	0.0125	0.025	0.015	0.0175	0.02	0.0225	0.005	0.0125
L2 weight decay for Inception	0.01	0.01	0.01	0.001	0.004	-	-	-
L1 weight decay for doctor models	0.001	0.00004	0.0001	0.001	0.01	-	-	-
L2 weight decay for doctor models	0.01	0.004	0.001	0.01	0.04	-	-	-
L1 weight decay for averaging logits	-	-	-	-	-	0.4	0.02	4
L2 weight decay for averaging logits	-	-	-	-	-	15	0.4	110
Bottleneck size	-	-	-	-	-	-	-	3

Table 3: Class distributions of training and validation datasets (as %).

Grade	Training	Validation
1	51.03	72.69
2	24.75	17.62
3	16.81	7.27
4	4.17	1.20
5	3.23	1.21

Table 4: Test metrics from Multi-class vs Binary loss for BN.

Test Metric (%)	Trained with binary loss	Trained with 5-class loss
Binary AUC	95.58	97.11
Binary Error	11.27	9.92
Spec@97%Sens	63.12	79.60

Table 5: Summary of Results. All models in this table were trained with 5-class loss except DN Mnih, whose loss was the negative log posterior.

Dataset	Metric (%)	BN		DN		WDN	IWDN	BIWDN
		∅	EM	∅	Mnih			
Test	5-class Error	23.83	23.74	21.86	22.76	**20.58**	20.63	20.83
	Binary AUC	97.11	97.00	97.28	97.42	**97.45**	97.43	97.41
	Binary Error	9.92	10.12	9.75	10.24	**9.07**	9.12	9.23
	Spec@97%Sens	79.60	79.97	81.81	**83.61**	82.69	82.46	82.46
Validation	5-class Error	9.71	9.83	9.41	9.59	**9.35**	9.40	9.42
	Binary AUC	99.10	**99.17**	99.07	98.95	99.09	99.06	99.07
	Binary Error	3.02	3.13	3.19	3.65	**2.76**	2.87	2.89
	Spec@97%Sens	**94.60**	93.85	93.13	92.52	93.53	93.19	93.42

5.2 Averaging modeled doctors beats modeling the average doctor

We saw a reduction in 5-class classification test error of 1.97% from 23.83% (8.27% relative decrease) due to averaging modeled doctors (DN) instead of modeling the averaged doctor (BN). In comparison, using labels calculated with EM to train BN only reduced 5-class test classification error by 0.09% (0.38% relative decrease). Over BN, DN also increased binary AUC by 0.17% from 97.11%, decreased binary classification error by 0.17% from 9.92%, and increased specificity at 97% sensitivity (spec@97%sens) by 2.21% from 79.60%. Meanwhile, using labels calculated with EM on BN merely increased spec@97%sens by 0.37% compared to vanilla BN and actually led to slightly worse performance on binary AUC (-0.11%) and binary error (+0.20%). Note that the binary AUC, binary error, and spec@97%sens metrics could have been improved for all models had we done hyperparameter tuning and early stopping for them specifically, but we decided to do all our model selection on one metric (5-class error) both for simplicity and to simulate the metric decision required in real-life automated diagnosis systems. We see that DN was significantly better on all test metrics compared to BN trained using the labels obtained with EM.

5.3 Learning averaging weights helps

We saw a further 1.28% decrease in 5-class test error relative to BN from using WDN as opposed to DN (5.37% additional relative decrease). Binary AUC increased an additional 0.17%, binary classification error decreased another 0.68%, and spec@97%sens increased an extra 0.88%, all on test data. Results from IWDN and BIWDN were slightly worse than those from WDN. We would expect a bigger improvement from WDN and potentially further improvements from training averaging logits in an image-specific way if we had doctors with more varied abilities and greater environmental differences, but for the dataset we used image-specific averaging logits did not help. Our extension of Mnih and Hinton [15]'s competing algorithm actually caused DN to perform worse by 0.90% on 5-class test error (3.78% less relative reduction), and was also more computationally costly than (B)(I)WDN. A different noise model we considered did not help either (Appendix C.3).

6 CONCLUSIONS

We introduce a method to make more effective use of noisy labels when every example is labeled by a subset of a larger pool of experts. Our method learns from the identity of multiple noisy annotators by modeling them individually with a shared neural net that has separate sets of outputs for each expert, and then learning averaging weights for combining their modeled predictions. We evaluate our method on the diagnosis of diabetic retinopathy severity on the 5-point scale from images of the retina. Compared to our baseline model of training on the average doctor opinion, a strategy that yielded state-of-the-art results on automated diagnosis of DR, our method can lower 5-class classification test error from 23.83% to 20.58%, a relative reduction of 13.6%. We also found that, on binary metrics, training with a 5-class loss significantly beats training with a binary loss, as was done in the published baseline. We compared our method to competing algorithms by Welinder and Perona and by Mnih and Hinton and we showed that corresponding parts of our method give superior performance to both. Our methodology is generally applicable to supervised training systems using datasets with labels from multiple annotators.

A CODE

The TensorFlow code used in this paper will soon be made publicly available.

B MUTUAL INFORMATION FOR NOISY LABELS

Here we compute the mutual information between a noisy MNIST label and the truth, assuming random noise, in order to estimate the number of noisily labeled training cases equivalent to one case that is known to be correctly labeled.

Empirically, N perfectly labeled training cases give about the same test error as $NI_{\text{perfect}}/I_{\text{noisy}}$ training cases with noisy labels, where I_{noisy} is the mutual information per case between a noisy label and the truth and I_{perfect} is the corresponding mutual information for perfect labels. For ten classes, the mutual information (in nats) is $I_{\text{perfect}} = 2.3 = -log(0.1)$, but when the noisy label is

20% correct on average, the mutual information is:

$$I_{\text{noisy}} = 0.044 = -log(0.1) - 10 \times 0.02 \times log\left(\frac{0.1}{0.02}\right) - 90 \times 0.1 \times \frac{0.8}{9} log\left(\frac{0.1}{0.1 \times 0.8/9}\right).$$

So if the learning is making good use of the mutual information in the noisy labels we can predict that 60,000 noisy labels are worth $60,000 \times 0.044/2.3 \approx 1,148$ clean labels. In reality we needed about 1,000 clean labels to get similar results.

C OTHER IDEAS TESTED

C.1 Mean Class Balancing

In addition to log prior correction of class distributions, we also attempted mean class balancing wherein examples from less frequent classes are upweighted and more frequent classes are downweighted in the cross entropy loss, in inverse proportion to their prevalence relative to the uniform distribution across classes. Explicitly, we weight each example of class c by:

$$\alpha_c = \frac{\bar{q}}{q(c)} = \frac{1}{|c|q(c)},$$

Eigen and Fergus. [5] employ a similar method for computer vision tasks although they use medians instead of means. In our case, using mean class balancing lowered performance, possibly because it made too many assumptions on the hidden test distribution, and was not employed.

C.2 Alternative target distribution for training averaging logits

To train the averaging logits, we took each training case and use the opinions of the doctors who actually labeled that case to define the target output distribution. Alternatively, the target distribution can be defined as the equally weighted average of the predictions of the doctor models corresponding to the doctors who labeled that case. In the notation used in Table 1, this would be $\frac{1}{|I|}\sum_{i \in I} p_i$. We experimented with using this alternative target distribution in calculating cross entropy loss but saw inferior results.

C.3 A alternative noise model

Because the multi-class extension of Mnih and Hinton [15] we tried showed poor results, which we postulated may have been because it was sensitive to differences in class distributions between datasets, we considered a different noise model that made less assumptions on the class distribution of the data. We assumed a symmetric noise distribution that is determined by a single prior parameter. This assumes that if a label is wrong, it has equal probability of belonging to any of the other classes. However we allowed this parameter to vary by doctor. For each doctor d we estimated this parameter:

$$\theta_d = p(\tilde{m}_d = l | m_d = l)$$

with the real doctor reliability score calculated from the Welinder and Perona [24] algorithm. This method performed slightly worse than the 5-class variant of Mnih and Hinton [15]. Note that a number of other noise models of varying complexity can be considered as well.

D HYPERPARAMETER TUNING

For the MNIST experiment we used default Adam optimizer hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 1{\times}10^{-8}$. We did a grid search on learning rates in the set {0.000003, 0.00001, 0.00003, . . ., 0.003} and standard deviations of the initial random normal weights in the set {0.0001, 0.0003, 0.001, . . ., 0.01} and found optimal values of 0.00003 for the former and 0.001 for the latter.

For computer-aided diagnosis of DR we did a grid search on the following hyperparameter spaces: dropout for Inception backbone $\in$ {0.5, 0.55, 0.6, . . ., 1.0}, dropout for doctor models $\in$ {0.5, 0.55, 0.6, . . ., 1.0}, learning rate $\in$ {$1{\times}10^{-7}$, $3{\times}10^{-7}$, $1{\times}10^{-6}$, . . ., 0.03}, entropy weight $\in$ {0.0, 0.0025, 0.005, . . ., 0.03} $\cup$ {0.1}, weight decay for Inception $\in$ {0.000004, 0.00001, 0.00004, . . ., 0.1}, L1 weight decay for doctor models $\in$ {0.000004, 0.00001, 0.00004, . . ., 0.04}, L2 weight decay for doctor models $\in$ {0.00001, 0.00004, . . ., 0.04}, L1 weight decay for averaging logits $\in$ {0.001, 0.01, 0.02, 0.03, . . ., 0.1, 0.2, 0.3, . . ., 1, 2, 3, . . ., 10, 100, 1000}, L2 weight decay for averaging logits $\in$ {0.001, 0.01, 0.1, 0.2, 0.3, . . .,1, 5, 10, 15, 20, 30, . . ., 150, 200, 300, 400, 500, 1000}, and bottleneck size (for BIWDN) $\in$ {2, 3, 4, 5, 6, 7}. We used a learning rate decay factor of 0.99 optimized for BN. The magnitudes of the image preprocessing perturbations were also tuned for BN.

E DATASET DETAILS

119,589 of our training set images are the same as those used in the training set of Gulshan et al. [10] (which consists of 128,175 images). The images removed from the training dataset used by Gulshan et al. [10] are detailed here: (i) 4,204 out of the 128,175 were removed to create a separate validation dataset for experiments within the research group. (ii) 4,265 out of the 128,175 images were excluded since they were deemed ungradable by every ophthalmologist that graded them. Unlike Gulshan et al. [10], we do not predict image gradeability in this work and hence exclude those images. (iii) 117 out of the 128, 175 fail our image scale normalization preprocessing step and were also excluded. We also acquired 6,933 more labeled images since the creation of the training dataset in Gulshan et al. [10] and added them to this training set.

The validation dataset consists of 7,963 images obtained from EyePACS clinics. These images are a random subset of the 9,963 images of the EyePACS-1 test set used in Gulshan et al. [10]. The remaining 2,000 images were included as part of the test set in this work. In practice, only 7,805 of the 7,963 validation images have at least one label, since the remaining 158 images were of poor quality and considered ungradable by all ophthalmologists that labeled them.

The test set consists of 1,748 images of the Messidor-2 dataset [4] and the remaining 2,000 out of the 9,963 images of the EyePACS-1 test dataset used in Gulshan et al. [10]. 1,803 of the 2,000 images from the EyePACS-1 test set, and 1,744 of the 1,748 images of the Messidor-2 were considered gradable after adjudication and were assigned labels.

F MORE DIFFERENCES FROM PUBLISHED BASELINE

Here we list distinctions between BN and the model in Gulshan et al. [10] that are not mentioned in Section 4.3. Gulshan et al. [10]

defined referable diabetic retinopathy as the presence of moderate and worse diabetic retinopathy *or* referable diabetic macular edema, while we ignore information on the latter. They also only reported binary (referable/non-referable) classification metrics while we reported both binary and 5-class classification metrics. Finally, we did not ensemble replicates as Gulshan et al. [10] did because we focused on comparing different methods of using the labels rather than squeezing the last drop of performance from one method.

ACKNOWLEDGMENTS

Thanks to Dale Webster, Lily Peng, Jonathan Krause, Arunachalam Narayanaswamy, Quoc Le, Alexey Kurakin, Anelia Angelova, Nathan Silberman, George Dahl, Brian Cheung, Anna Goldie, David Ha, Matt Hoffman, Olga Wichrowska, Justin Gilmer, Denny Britz, Mohammad Norouzi, and Luke Metz for helpful discussions and feedback. This work was supported by the Google Brain Residency program, for whom we give particular thanks to Leslie Phillips, Samy Bengio, and Jeff Dean.

REFERENCES

[1] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. 2016. TensorFlow: A System for Large-Scale Machine Learning. *OSDI* (November 2016).

[2] M. D. Abràmoff, J. C. Folk, D. P. Han, J. D. Walker, D. F. Williams, S. R. Russell, P. Massin, B. Cochener, L. Tang P. Gain, M. Lamard, D. C. Moga, G. Quellec, and M. Niemeijer. 2013. Automated analysis of retinal images for detection of referable diabetic retinopathy. *JAMA Ophthalmol.* 131, 3 (March 2013), 351–357.

[3] A. P. Dawid and A. M. Skene. 1979. Maximum likelihood estimation of observer error-rates using the em algorithm. *Applied Statistics* 28, 1 (March 1979), 20–28.

[4] E. Decencière, Z. Xiwei, G. Cazuguel, B. Lay, B. Cochener, C. Trone, P. Gain, J.-R. Ordó nez Varela, P. Massin, A. Erginay, B. Charton, and J.-C. Kleain. 2014. Feedback on a Publicly Distributed Image Database: The Messidor Database. *Image Anal Stereol* 33, 3 (July 2014), 231–234.

[5] D. Eigen and R. Fergus. 2015. Predicting Depth, Surface Normals and Semantic Labels with a Common Multi-Scale Convolutional Architecture. *ICCV* 11, 18 (December 2015), 2650–2658.

[6] J. G. Elmore, G. M. Longton, P. A. Carney, B. M. Geller, T. Onega, A .N. A. Tosteson, H. D. Nelson, M. S. Pepe, K. H. Allison, S. J. Schnitt, F. P. O'Malley, and D. L. Weaver. 2015. Diagnostic concordance among pathologists interpreting breast biopsy specimens. *JAMA* 313, 11 (March 2015), 1122–1132.

[7] J. G. Elmore, C. K. Wells, C. H. Lee, D. H. Howard, and A. R. Feinstein. 1994. Variability in radiologists' interpretations of mammograms. *NEJM* 331, 22 (December 1994), 1493–1499.

[8] International Diabetes Foundation. 2015. IDF Diabetes Atlas, 7th Ediction. (2015). diabetesatlas.org.

[9] I. Goodfellow, Y. Bengio, and A. Courville. 2016. *Deep Learning*. MIT Press, Cambridge, Massachusetts.

[10] V. Gulshan, L. Peng, M. Coram, M. C. Stumpe, A. Narayanaswamy D. Wu, S. Venugopalan, K. Widner, T, Madams, J. Cuadros, R. Kim, R. Raman, P. C. Nelson, J. L. Mega, and D. R. Webster. 2016. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA* 316, 22 (December 2016), 2402–2410.

[11] Kaggle Inc. 2015. Kaggle Diabetic Retinopathy Detection competition. (July 2015). kaggle.com/c/diabetic-retinopathy-detection.

[12] National Eye Institute. 2015. Facts About Diabetic Eye Disease. (September 2015). nei.nih.gov/health/diabetic/retinopathy.

[13] S. Ioffe and C. Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ICML* 37 (March 2015), 448–45.

[14] D. P. Kingma and J. L. Ba. 2015. ADAM: A Method for Stochastic Optimization. *ICLR* (July 2015).

[15] V. Mnih and G. E. Hinton. 2012. Learning to Label Aerial Images from Noisy Data. *ICML* (July 2012), 567–574.

[16] American Academy of Ophthalmology. 2002. International clinical diabetic retinopathy disease severity scale. (October 2002). Detailed table. icoph.org/dynamic/attachments/resources/diabetic-retinopathy-detail.pdf.

[17] G. Pereyra, G. Tucker, L. Kaiser, and G. E. Hinton. 2017. Regularizing Neural Networks by Penalizing Confident Output Distributions. *arXiv* (January 2017).

[18] V. Raykar, S. Yu, L. Zhao, A. Jerebko, C. Florin, G. Valadez, L. Bogoni, and L. Moy. 2009. Supervised Learning from Multiple Experts: Whom to trust when everyone lies a bit. *ICML* 26 (June 2009), 889—-896.

[19] V. C. Raykar and S. Yu. 2012. Eliminating spammers and ranking annotators for crowdsourced labeling tasks. *JMLR* 13 (2012), 491—-518.

[20] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *JMLR* 115, 3 (January 2015), 211–252.

[21] P. Smyth, U. Fayyad, M. Burl, P. Perona, and P. Baldi. 1995. Inferring ground truth from subjective labelling of Venus images. *NIPS* 7 (1995), 1085–92.

[22] K. Solanki, C. Ramachandra, S. Bhat, M. Bhaskaranand, M. G. Nittala, and S. R. Sadda. 2015. EyeArt: Automated, High-throughput, Image Analysis for Diabetic Retinopathy Screening. *Invest Ophthalmol Vis Sci* 56, 7 (June 2015). Meeting abstract.

[23] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. 2016. Re-thinking the inception architecture for computer vision. *CVPR* (June 2016), 2818–2826.

[24] P. Welinder and P. Perona. 2010. Online crowdsourcing: rating annotators and obtaining cost-effective labels. *CVPR Workshop* (June 2010), 25–32.

[25] J. Whitehill, P. Ruvolo, T. Wu, J. Bergsma, and J. Movellan. 2009. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. *NIPS* 22 (2009), 2035—-2043.